

Difference Vector Equalization for Robust Fine-tuning of Vision-Language Models

Satoshi Suzuki, Shin’ya Yamaguchi, Shoichiro Takeda, Taiga Yamane, Naoki Makishima, Naotaka Kawata, Mana Ithori, Tomohiro Tanaka, Shota Orihashi, Ryo Masumura

NTT, Inc., Japan
satoshi xv.suzuki@ntt.com

Abstract

Contrastive pre-trained vision-language models, such as CLIP, demonstrate strong generalization abilities in zero-shot classification by leveraging embeddings extracted from image and text encoders. This paper aims to robustly fine-tune these vision-language models on in-distribution (ID) data without compromising their generalization abilities in out-of-distribution (OOD) and zero-shot settings. Current robust fine-tuning methods tackle this challenge by reusing contrastive learning, which was used in pre-training, for fine-tuning. However, we found that these methods distort the geometric structure of the embeddings, which plays a crucial role in the generalization of vision-language models, resulting in limited OOD and zero-shot performance. To address this, we propose **Difference Vector Equalization (DiVE)**, which preserves the geometric structure during fine-tuning. The idea behind DiVE is to constrain *difference vectors*, each of which is obtained by subtracting the embeddings extracted from the pre-trained and fine-tuning models for the same data sample. By constraining the difference vectors to be equal across various data samples, we effectively preserve the geometric structure. Therefore, we introduce two losses: average vector loss (AVL) and pairwise vector loss (PVL). AVL preserves the geometric structure globally by constraining difference vectors to be equal to their weighted average. PVL preserves the geometric structure locally by ensuring a consistent multimodal alignment. Our experiments demonstrate that DiVE effectively preserves the geometric structure, achieving strong results across ID, OOD, and zero-shot metrics.

Extended version — <https://arxiv.org/abs/2511.09973>

1 Introduction

Advances in contrastive pre-training have led to the development of vision-language models, such as CLIP (Radford et al. 2021) and ALIGN (Jia et al. 2021). These models consist of two separate encoders for visual and textual modalities. They are pre-trained on large-scale image-text datasets to learn multimodal representations by mapping images and texts into a shared embedding space. After pre-training, the model can execute classification by using these embeddings even for previously unseen tasks. Specifically, it prepares text prompts for possible class names and selects the class,

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

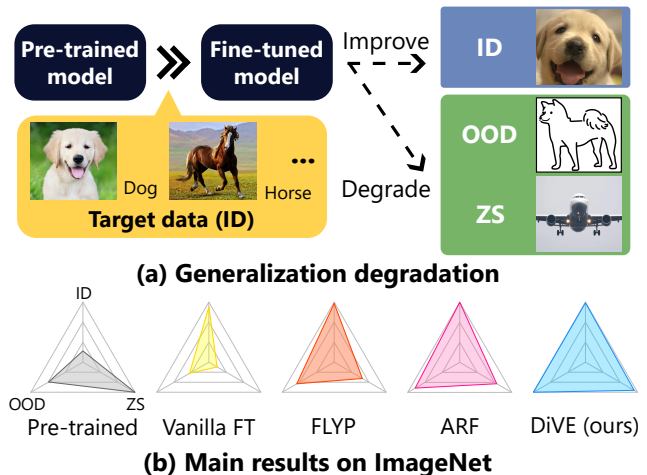


Figure 1: (a) Illustrative example of generalization degradation. Performance of pre-trained models in out-of-distribution (OOD) and zero-shot (ZS) settings severely degrades after vanilla fine-tuning (FT) on in-distribution (ID) data. (b) Normalized performance of robust fine-tuning methods on ID, OOD, and ZS metrics. We used ImageNet as target task. DiVE performs well across all metrics.

the text embedding of which has the highest similarity to the image embedding. Such “zero-shot” classification using vision-language models achieves reasonable performance on many tasks (Radford et al. 2021) and impressive robustness to various distribution shifts, such as those from natural images to sketches (Wang et al. 2019).

Despite advances in vision-language models, their zero-shot performance is often sub-optimal, especially on data samples that deviate from the pre-training domain, such as satellite images (Radford et al. 2021). In such cases, task-specific fine-tuning is commonly used to enhance performance. However, fine-tuning a pre-trained model can easily degrade its generalization ability. Specifically, vanilla fine-tuning (FT) improves performance on in-distribution (ID) target data but significantly degrades out-of-distribution (OOD) and zero-shot performance (Kumar et al. 2022; Ilharco et al. 2022). Figure 1 (a) illustrates an example: while vanilla FT improves classification perfor-

mance on ID data samples (e.g., dog), it degrades OOD performance on samples from the same classes but in different styles (e.g., sketch), as well as zero-shot performance on samples from classes not seen during fine-tuning (e.g., airplane). Given this phenomenon, our goal is to maximize ID performance without compromising generalization ability. Achieving this could enhance the practical value of vision-language models in diverse and dynamic environments such as self-driving cars.

To tackle this challenge, robust fine-tuning methods for vision-language models have been proposed. For example, FLYP (Goyal et al. 2023) executes contrastive learning during fine-tuning using class label prompts (e.g., “a photo of a dog”), unlike vanilla FT, which updates only the image encoder (Wortsman et al. 2022). This strategy is expected to enhance the generalization ability by considering both image and text encoders. Subsequently, ARF (Han et al. 2024) introduces an auxiliary contrastive loss into FLYP using a reference image-caption dataset with diverse semantics. This approach works as a replay technique (Rolnick et al. 2019), retaining pre-trained knowledge.

While these robust fine-tuning methods demonstrate promising performance, there may still be room for improvement in both OOD and zero-shot performance. As one interesting fact, we found that these methods distort the geometric structure of the embeddings learned by pre-training (see Table 1), i.e., the relative positions between the embeddings. Since the relative distances between embeddings reflect the semantic similarity of corresponding inputs, this structure is fundamental to the strong generalization ability of vision-language models (Goel et al. 2022; Oh et al. 2023; Yamaguchi et al. 2025). Consequently, its distortion can degrade both OOD and zero-shot performance. Therefore, the following research question naturally arises: *Can we improve robust fine-tuning for vision-language models by preserving the geometric structure of the embeddings?*

In this paper, we propose a novel robust fine-tuning method **Difference Vector Equalization** (DiVE). DiVE builds on FLYP, i.e., using contrastive loss for fine-tuning, and introduces a constraint to preserve the geometric structure of the embeddings during fine-tuning on the target data. The idea behind DiVE is to constrain *difference vectors*, each of which is obtained by subtracting the embeddings extracted from the pre-trained and fine-tuning models for the same data sample. Intuitively, difference vectors represent the changes in the samples in the embedding space caused by fine-tuning; thus, constraining difference vectors to be equal across various data samples can help preserve the geometric structure (see Fig. 2). Therefore, we first prepare the image-caption reference dataset, which is also used with ARF, to obtain the difference vectors for various images and captions. We then introduce two new losses for DiVE: *average vector loss* (AVL) and *pairwise vector loss* (PVL). AVL computes a weighted average of the difference vectors and constrains each difference vector to be equal to this average vector. Therefore, it preserves the global geometric structure formed by various image and text embeddings. On the other hand, PVL constrains the difference vector for each reference image to be equal to that for its corresponding caption.

It preserves the local geometric structure by ensuring consistent image-caption alignment for each pair. By incorporating AVL and PVL into FLYP, DiVE can effectively preserve the geometric structure of embeddings during fine-tuning.

We empirically demonstrate that DiVE effectively preserves the geometric structure of the embeddings learned by pre-training, as expected (see Table 1). Concurrently, DiVE performs impressively on various classification tasks. For example, Fig. 1 (b) shows that it consistently performs well across ID, OOD, and zero-shot metrics. These results highlight the importance of preserving the geometric structure for robust fine-tuning.

2 Related Work

Robust Fine-tuning for vision-language models. Vision-language models, such as CLIP (Radford et al. 2021), demonstrate strong generalization abilities in OOD and zero-shot settings. However, empirical evidence suggests that vanilla FT, which updates the weights of the image encoder using downstream data, compromises these strengths (Wortsman et al. 2022; Ilharco et al. 2022). To tackle this challenge, robust fine-tuning methods for vision-language models have been proposed. An early baseline method, LP-FT (Kumar et al. 2022), first applies linear probing then fine-tunes the entire weights to enhance OOD performance. FLYP (Goyal et al. 2023) uses vision-language contrastive learning for fine-tuning. This choice can enhance OOD performance over vanilla FT by considering both image and text encoders. ARF (Han et al. 2024) is the first robust fine-tuning method that explicitly aims to enhance not only OOD but also zero-shot performance. It combines FLYP and a replay technique (Rolnick et al. 2019), which involves training on a reference dataset (e.g., CC3M (Sharma et al. 2018)) that resembles the pre-training data, to help recover the inherent generalization ability of the pre-trained model. While FLYP and ARF show promising performance, both OOD and zero-shot performance can still be improved. DiVE improves them by introducing a new perspective: preserving the geometric structure of the embeddings.

Geometry of Embeddings of Contrastive Learning Models. Contrastive learning has been widely used for representation learning on various modalities (Chen et al. 2020; Gao, Yao, and Chen 2021). Most contrastive learning methods adopt L_2 -normalized embeddings for stable learning (Xu and Durrett 2018), which causes the embeddings to lie on a unit hypersphere. Wang and Isola (2020) found that the relative positions between hyperspherical embeddings, i.e., the geometric structure, play a crucial role in generalization in contrastive learning models. Studies (Goel et al. 2022; Oh et al. 2023; Yamaguchi et al. 2025) have shown that this structure is also crucial for contrastive pre-trained vision-language models, such as CLIP. Among these, CyCLIP (Goel et al. 2022) shares some concepts with DiVE in constraining the geometric structure of the embeddings. Specifically, it constrains the cosine similarities between in-modal and cross-modal embeddings during pre-training. We found that such a cosine similarity-based constraint is insufficient to preserve the geometric structure and improve

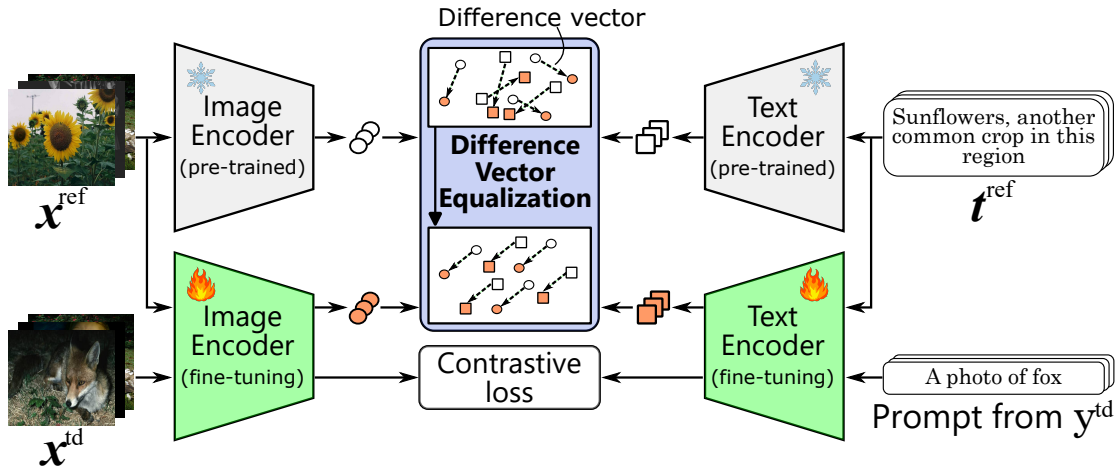


Figure 2: Overview of our proposed method, **Difference Vector Equalization (DiVE)**. It uses contrastive loss for fine-tuning. While fine-tuning on target data, it constrains all difference vectors for reference data (x^{ref} and t^{ref}) to be equal.

performance, as it captures only the angular relationship between embeddings. In contrast, our difference vector-based constraint achieves both (see Table 7).

Maintaining Zero-shot Performance. There are several approaches for maintaining the zero-shot performance of vision-language models during fine-tuning, not only robust fine-tuning. Prompt learning (Zhou et al. 2022b,a; Khatkhat et al. 2023) optimizes a set of learnable prompt vectors, while freezing the weights of the pre-trained model. This strategy perfectly maintains the generalization ability. However, as reported by (Shu et al. 2023), its classification performance on target data is often limited due to the small number of learnable parameters. Methods for continual learning (Wang et al. 2024) can be adapted to our setting by interpreting the contrastive pre-training as the previous task. For example, the state-of-the-art method SnD (Yu et al. 2024) is related to DiVE. Specifically, it leverages a reference dataset (e.g., CC3M) for feature distillation, i.e., constraining the image embeddings from the pre-trained and fine-tuning models to be identical during continual learning. In our context, this can be viewed as forcing the difference vectors to be zero. However, such a strong constraint may hinder adaptation to target data. In contrast, DiVE allows for non-zero difference vectors, offering greater flexibility and better performance (see Tables 2, 3, and 4).

Weight Ensemble. In an orthogonal line of research, several studies (Wortsman et al. 2022; Ilharco et al. 2022) demonstrated that ensembling the weights of pre-trained and fine-tuned models can improve OOD and zero-shot performance. Although they were originally proposed in integration with vanilla FT, DiVE can also benefit from such ensembling strategies. We present the results of integrating DiVE into this strategy in the Appendix.

3 Preliminaries

We fine-tune a contrastive pre-trained vision-language model, such as CLIP (Radford et al. 2021), on target data. A vision-language model comprises an image encoder f_θ :

$\mathcal{X} \rightarrow \mathbb{R}^d$ and text encoder $g_\phi : \mathcal{T} \rightarrow \mathbb{R}^d$, where θ and ϕ represent the weights of the image and text encoders, respectively. Furthermore, \mathcal{X} and \mathcal{T} denote the image and text spaces, respectively, and d is the dimension of the embeddings. During the pre-training, image-text pairs are sampled from a web-scale training dataset $\mathcal{S} = \{(x_i, t_i)\}_{i=1}^N$, where $x_i \in \mathcal{X}$ is the i -th image sample and $t_i \in \mathcal{T}$ is the corresponding text description. A contrastive loss function \mathcal{L}_{cl} is used to align the image embedding $f_\theta(x_i)$ with the corresponding text embedding $g_\phi(t_i)$,

$$\mathcal{L}_{\text{cl}} = -\frac{1}{2B} \sum_{i=1}^B \log \frac{\exp(f_\theta(x_i) \cdot g_\phi(t_i)/\tau)}{\sum_{k=1}^B \exp(f_\theta(x_i) \cdot g_\phi(t_k)/\tau)} - \frac{1}{2B} \sum_{i=1}^B \log \frac{\exp(f_\theta(x_i) \cdot g_\phi(t_i)/\tau)}{\sum_{k=1}^B \exp(f_\theta(x_k) \cdot g_\phi(t_i)/\tau)}. \quad (1)$$

Here, B is the batch size, τ denotes the temperature used to scale the pairwise similarities, and \cdot indicates an inner product. Both encoders output L_2 -normalized embeddings, making their inner product equivalent to cosine similarity.

During inference, we first prepare text prompts $t^c \in \mathcal{T}$ for each class of interest $c \in \mathcal{C}$, such as

$$t^c = \text{“a photo of a [c]”}, \quad (2)$$

where $[c]$ represents the c -th class name. Even for previously unseen tasks, i.e., the zero-shot setting, the model classifies a test image x^{test} by selecting the class \hat{y} whose t^c has the highest similarity to the image embedding,

$$\hat{y} = \arg \max_{c \in \mathcal{C}} f_\theta(x^{\text{test}}) \cdot g_\phi(t^c). \quad (3)$$

We consider a fine-tuning problem that updates the weights of encoders, i.e., θ and ϕ , using the target data $\mathcal{S}^{\text{td}} = \{(x_i^{\text{td}}, y_i^{\text{td}})\}_{i=1}^{N_{\text{td}}}$. With DiVE, we inherit the basic fine-tuning strategy of FLYP, which leverages supervision via a contrastive loss. Namely, we use the loss function \mathcal{L}_{cl} in Eq. (1) on \mathcal{S}^{td} , where the class label y_i^{td} is converted into a text prompt, as described in Eq. (2).

4 Proposed Method

This paper proposes DiVE, which explicitly preserves the geometric structure of the embeddings. Figure 2 shows an overview of DiVE. Similar to ARF (Han et al. 2024), we use a reference image-caption dataset $\mathcal{S}^{\text{ref}} = \{(\mathbf{x}_i^{\text{ref}}, \mathbf{t}_i^{\text{ref}})\}_{i=1}^{N_{\text{ref}}}$, which resembles the pre-training dataset \mathcal{S} and covers diverse semantics. Our main idea is to constrain the difference vectors between the embeddings extracted from the pre-trained and fine-tuning models to be equal across all images and captions in \mathcal{S}^{ref} . The difference vectors for $\mathbf{x}_i^{\text{ref}}$ and $\mathbf{t}_i^{\text{ref}}$ are defined as¹,

$$u(\mathbf{x}_i^{\text{ref}}) = f_{\theta_{\text{ft}}}(\mathbf{x}_i^{\text{ref}}) - f_{\theta_{\text{pre}}}(\mathbf{x}_i^{\text{ref}}), \quad (4)$$

$$v(\mathbf{t}_i^{\text{ref}}) = g_{\phi_{\text{ft}}}(\mathbf{t}_i^{\text{ref}}) - g_{\phi_{\text{pre}}}(\mathbf{t}_i^{\text{ref}}), \quad (5)$$

where $\{\theta_{\text{pre}}, \phi_{\text{pre}}\}$ and $\{\theta_{\text{ft}}, \phi_{\text{ft}}\}$ represent the weights of the pre-trained and fine-tuning models, respectively. In other words, our ideal state is $u(\mathbf{x}_1^{\text{ref}}) = v(\mathbf{t}_1^{\text{ref}}) = \dots = u(\mathbf{x}_{N_{\text{ref}}}^{\text{ref}}) = v(\mathbf{t}_{N_{\text{ref}}}^{\text{ref}})$ after fine-tuning. Achieving this state effectively preserves the geometric structure of the embeddings extracted from the pre-trained model (see Fig. 2), protecting it from distortion caused by fine-tuning (Rajaei and Pilehvar 2021; Kumar et al. 2022). We thus introduce two new losses, i.e., average vector loss (AVL) and pairwise vector loss (PVL), which we formulate in the following sections.

4.1 Average Vector Loss

In AVL, we first compute an average vector $\mathbf{m} \in \mathbb{R}^d$ from the difference vectors for the images and captions within a mini-batch from \mathcal{S}^{ref} . Here, \mathbf{m} is defined as the exponential moving average of difference vectors from previous batches,

$$\mathbf{m} = \alpha \cdot \mathbf{m}_p + \frac{(1 - \alpha)}{B'} \sum_{j=1}^{B'} \left(\frac{u(\mathbf{x}_j^{\text{ref}}) + v(\mathbf{t}_j^{\text{ref}})}{2} \right), \quad (6)$$

where \mathbf{m}_p is the average vector used in the immediately previous batch, and B' is the batch size of the mini-batch from \mathcal{S}^{ref} . B' does not necessarily have to be equal to the batch size of the mini-batch from target data (i.e., B in Eq. (1)). Furthermore, α is a hyper-parameter for determining the effect of the current difference vectors on updating the average vector. By setting a large α (e.g., 0.99), the average vector maintains high consistency across batches while reflecting the current trend of difference vectors. We initialize \mathbf{m} as a zero vector because \mathbf{m}_p is undefined for the first batch.

By using \mathbf{m} , AVL constrains each difference vector,

$$\mathcal{L}_{\text{avl}} = \frac{1}{B'} \sum_{j=1}^{B'} (\|u(\mathbf{x}_j^{\text{ref}}) - \mathbf{m}\|^2 + \|v(\mathbf{t}_j^{\text{ref}}) - \mathbf{m}\|^2). \quad (7)$$

This formulation encourages all difference vectors to be equal to the average vector and preserves the global geometric structure formed by various image and text embeddings.

¹Strictly, embeddings lie on a unit hypersphere, so difference vectors should account for curvature. However, we found that the norm of the difference vectors is very small (10^{-3} to 10^{-4} on average), allowing the curvature to be ignored in practice (Lee 2012).

Method	RSA correlation score (\uparrow)
FLYP	0.825
FLYP + replay	0.850
SnD	0.847
FLYP + AVL	0.978
FLYP + PVL	0.976
DiVE	0.981

Table 1: Representation similarity analysis (RSA) correlation scores of robust fine-tuning methods. ‘‘FLYP + replay’’ denotes FLYP using auxiliary contrastive loss, which serves as proxy for ARF. We used Flickr8K dataset for evaluation.

4.2 Pairwise Vector Loss

While AVL preserves the global geometric structure, PVL preserves the local geometric structure formed by each image-caption pair. Since distortion of this local structure directly affects inference of the vision-language model using Eq. (3), preserving it is critical for their generalization. PVL thus constrains the difference vectors for the reference image $\mathbf{x}_j^{\text{ref}}$ and that for its caption $\mathbf{t}_j^{\text{ref}}$,

$$\mathcal{L}_{\text{pvl}} = \frac{1}{B'} \sum_{j=1}^{B'} \|u(\mathbf{x}_j^{\text{ref}}) - v(\mathbf{t}_j^{\text{ref}})\|^2. \quad (8)$$

Finally, by combining AVL and PVL into FLYP, we obtain the loss for fine-tuning the pre-trained model,

$$\mathcal{L}_{\text{final}} = \mathcal{L}_{\text{cl}} + \lambda \cdot (\mathcal{L}_{\text{avl}} + \mathcal{L}_{\text{pvl}}), \quad (9)$$

where λ is the hyper-parameter for determining the effect of AVL and PVL. Furthermore, \mathcal{L}_{cl} is the contrastive loss defined in Eq. (1) using the target data \mathcal{S}^{td} . Note that during inference, DiVE uses only the final fine-tuned model, incurring no additional computational cost beyond the original vision-language models.

5 Experiments

5.1 Datasets

We used the following benchmarks to measure ID, OOD, and zero-shot performance.

ImageNet (Deng et al. 2009) is an object-centric dataset containing 1.2M training and 50,000 validation images from 1,000 classes. We considered ImageNet-V2 (Recht et al. 2019), ImageNet-R (Hendrycks et al. 2021a), ImageNet-A (Hendrycks et al. 2021b), ImageNet-Sketch (Wang et al. 2019), and ObjectNet (Barbu et al. 2019) as the OOD datasets following (Goyal et al. 2023; Han et al. 2024). Following common practice (He et al. 2016), we report the performance on the ImageNet validation set as ID performance. **iWildCam** (Beery, Cole, and Gjoka 2020) consists of images from 182 animal species. The ID and OOD datasets differ in the camera used and factors such as background and illumination. We used the standard splits for training, validation, and ID and OOD testing (Koh et al. 2021). As the dataset has label imbalance, we report the macro F1-score, following Goyal et al. (2023).

Method	ImageNet (ID)	ImageNet-V2	ImageNet-R	ImageNet-A	ImageNet-Sketch	ObjectNet	OOD avg.
Pre-trained	68.3	61.9	77.7	50.0	48.3	55.4	58.7
Vanilla FT	81.3	71.2	66.1	37.8	46.1	53.3	54.9
LP-FT	81.7	72.1	73.5	47.6	50.3	58.2	60.3
FLYP [†]	82.2	73.0	71.5	48.4	49.7	54.8	59.5
ARF	82.7	72.8	75.6	50.3	51.8	55.8	61.3
SnD [†]	82.4	73.2	74.3	50.0	51.4	54.5	60.7
DiVE (ours)	82.5	73.8	77.3	54.9	52.9	56.9	63.2

Table 2: Experimental results of ID and OOD performance using CLIP ViT-B/16. We report results of ID dataset, i.e., ImageNet, and five OOD datasets. Numbers represent top-1 accuracy. [†] denotes our reproduction results. Bold indicates best result among fine-tuned methods, i.e., excluding pre-trained model.

FMoW (Christie et al. 2018) consists of satellite images of 62 classes for different land and building types. The ID and OOD datasets differ in the time of their collection and location, i.e., continent. We used the standard splits for training, validation, and ID and OOD testing (Koh et al. 2021). We report the ID testing accuracy and the worst-region OOD testing accuracy, following Goyal et al. (2023).

Zero-shot performance is evaluated using ten image classification datasets that cover various image domains: Caltech-101 (Fei-Fei, Fergus, and Perona 2004), Flowers (Nilsback and Zisserman 2008), Food (Bossard, Guillaumin, and Van Gool 2014), SUN397 (Xiao et al. 2010), DTD (Cimpoi et al. 2014), Aircraft (Maji et al. 2013), Stanford-Cars (Krause et al. 2013), OxfordPets (Parkhi et al. 2012), EuroSAT (Helber et al. 2019), and UCF-101 (Soomro, Zamir, and Shah 2012). We followed the zero-shot evaluation strategy by (Radford et al. 2021) using prompt engineering.

We used CC3M (Sharma et al. 2018) as the reference image-caption dataset \mathcal{S}^{ref} to ensure fair comparison with ARF, which also uses this dataset. We also used Flickr8K (Rashtchian et al. 2010) and COCO Captions (Lin et al. 2014) for the analysis in Sec. 6.

5.2 Baselines and Implementation Details

We considered three key baselines: FLYP (Goyal et al. 2023), ARF (Han et al. 2024), and SnD (Yu et al. 2024). The first two are robust fine-tuning methods that show promising performance, as discussed in Sec. 2. SnD was originally designed for continual learning, but we adapted it to our setting. Specifically, during fine-tuning, it constrains image embeddings as $f_{\theta_{\text{ft}}}(\mathbf{x}_j^{\text{ref}}) = f_{\theta_{\text{pre}}}(\mathbf{x}_j^{\text{ref}})$ using a reference dataset. We also report the performance of the pre-trained model, vanilla FT (Wortsman et al. 2022), and LP-FT (Kumar et al. 2022) for a comprehensive comparison. Both vanilla FT and LP-FT update only the image encoder using the cross-entropy loss with an added classification head.

We used CLIP ViT-B/16 (Dosovitskiy et al. 2021) pre-trained on LAION-400M (Schuhmann et al. 2021) as the base pre-trained model and also used CLIP ViT-L/14 in Table 8. We report the average results across three seeds. We reused the codebase and most of the hyper-parameters from FLYP (Goyal et al. 2023), including the AdamW optimizer (Loshchilov and Hutter 2019), cosine learning rate scheduler, and a batch size of $B = 512$ for ImageNet and $B = 256$ for other datasets. Unless otherwise noted, we set

$B' = B$. We also carried out early stopping on the basis of the ID validation performance. The Appendix describes the hyper-parameters and their selection in detail.

Regarding the hyper-parameters in DiVE, we set $\lambda = 1000$ chosen from $\{100, 500, 1000, 2500, 5000\}$ by tuning on the ImageNet validation set and used it for all datasets. We set $\alpha = 0.99$ in Eq. (6); its effect is analyzed in Sec. 5.5.

5.3 Effect on Preserving Geometric Structure

We first evaluated the effect of DiVE and the key baselines on preserving the geometric structure of the embeddings, a key aspect for our research question. We used the representation similarity analysis (RSA) correlation score (Kriegeskorte, Mur, and Bandettini 2008; Dwivedi and Roig 2019). This score quantifies the similarity between all pairwise relationships of embeddings obtained from two different models. Specifically, we computed the Pearson correlation between the pairwise cosine distance matrices of embeddings from the pre-trained model and from each fine-tuned model, using the Flickr8K dataset. A higher score indicates better preservation of the geometric structure of the embeddings extracted from the pre-trained model.

Table 1 lists the RSA correlation scores obtained with the baselines and DiVE, including its ablated variants, when using ImageNet as a target task. Since the official implementation of ARF is not publicly available, we used FLYP with an auxiliary contrastive loss on CC3M as a proxy for ARF, which we refer to as “FLYP + replay”. Its implementation details are explained in the Appendix. We found that the baseline methods struggled to fully preserve the geometric structure of the embeddings extracted from the pre-trained model, as indicated by the RSA correlation scores remaining in 0.825 to 0.850. In contrast, DiVE and its variants (with AVL or PVL only) significantly improved the preservation, achieving scores of about 0.98. These results indicate that DiVE successfully preserves the geometric structure, as intended. DiVE also achieved a higher score than either AVL or PVL alone, demonstrating the effectiveness of combining these complementary losses.

5.4 Performance Comparison with Baselines

We evaluated DiVE on ImageNet as a target task and compared its performance with the baselines. Table 2 reports ID and OOD performance, while Table 3 presents zero-shot performance. From Table 2, DiVE achieved the second-best

Method	Caltech	Flowers	Food	SUN	DTD	Aircraft	Cars	Pets	EuroSAT	UCF	ZS avg.
Pre-trained	88.0	71.0	88.5	65.3	44.7	24.3	64.7	89.0	55.3	68.7	66.0
Vanilla FT	78.8	16.0	37.3	39.3	29.7	4.7	10.8	80.2	15.4	44.3	35.7
LP-FT	84.0	44.3	68.8	49.9	37.9	15.8	37.7	81.9	30.4	59.5	51.0
FLYP [†]	87.6	39.7	63.3	52.6	36.8	8.0	32.3	77.2	38.2	59.0	49.5
ARF	88.6	46.4	74.5	63.8	40.4	13.9	44.7	83.1	35.8	64.6	55.6
SnD [†]	89.1	49.5	69.6	58.7	38.7	11.0	42.5	79.6	42.7	62.6	54.4
DiVE (ours)	88.4	66.0	84.3	64.7	47.0	22.1	55.5	88.4	51.4	68.9	63.7

Table 3: Zero-shot performance on ten evaluation datasets using CLIP ViT-B/16 on ImageNet. Numbers represent top-1 accuracy. [†] denotes our reproduction results. Bold indicates best result among fine-tuned methods, i.e., excluding pre-trained model.

Method	iWildCam			FMoW		
	ID	OOD	ZS	ID	OOD	ZS
Pre-trained	8.7	11.2	66.0	20.4	18.7	66.0
Vanilla FT	48.1	35.0	50.2	68.5	39.2	44.6
LP-FT	49.7	34.7	51.6	68.4	40.4	57.9
FLYP	52.2	35.6	51.0	68.6	41.3	45.1
FLYP + replay [†]	48.5	35.8	62.3	68.7	41.2	63.0
SnD [†]	50.6	37.0	60.4	67.0	41.4	56.6
DiVE	53.1	37.2	65.3	69.9	42.3	65.1

Table 4: Results on iWildCam and FMoW. ZS indicates average zero-shot performance across ten evaluation datasets. Following common practice, we report macro F1-score for iWildCam and top-1 accuracy for FMoW. [†] denotes our reproduction results. Bold indicates best result among fine-tuned methods, i.e., excluding pre-trained model.

ID performance, despite introducing additional constraints, AVL and PVL, into FLYP. For OOD performance, DiVE achieved state-of-the-art results on four out of five OOD datasets, as well as the highest average performance across them. For zero-shot performance in Table 3, DiVE achieved state-of-the-art performance on nine out of ten evaluation datasets, with an average performance of 63.7%. This represents a significant improvement over the previous best method, ARF, which achieved 55.6%. It is worth noting that the zero-shot performance dropped by only 2.3 points from the pre-trained CLIP ViT-B/16 while improving by 14.2 and 4.5 points in ID and OOD performance, respectively. As a result, DiVE achieved a well-balanced performance across the three evaluation metrics (ID, OOD, and zero-shot) at a high level, as shown in Fig. 1 (b). Given that DiVE significantly improves the preservation of the geometric structure of the embeddings learned by pre-training, as shown in Table 1, these results support that preserving the geometric structure is effective for robust fine-tuning, providing a positive answer to our research question. We also confirmed statistical significance on the basis of the results from three seeds, as reported in in the Appendix.

Table 4 also shows the results when iWildCam and FMoW are each used as a target task. We report the average zero-shot performance over ten evaluation datasets. Since ARF’s performance on iWildCam and FMoW is not reported, we used “FLYP + replay” as the proxy for ARF. Consistent with the results on ImageNet, DiVE demonstrated state-of-the-art performance across the three evaluation met-

Method		ImageNet		ZS avg.
AVL	PVL	ID	OOD avg.	
Baseline (FLYP)		82.2	59.5	49.5
✓		82.4	62.9 (+3.4)	62.9 (+13.4)
	✓	82.4	62.6 (+3.1)	62.7 (+13.2)
✓	✓	82.5	63.2 (+3.7)	63.7 (+14.2)

Table 5: Ablation study of AVL and PVL in DiVE on ImageNet. Baseline uses only contrastive loss (i.e., FLYP).

α	ID	OOD avg.	ZS avg.
0	82.4	61.7	61.4
0.5	82.5	62.6	62.4
0.9	82.4	62.8	62.9
0.99 (default)	82.4	62.9	62.9

Table 6: Effect of α of AVL on performance.

rics. These results indicate that the effectiveness of DiVE is not specific to a particular dataset.

As described in Sec. 2, SnD shares certain concepts with DiVE in that it constrains the difference vectors to be zero. However, from Tables 2, 3, and 4, SnD consistently underperformed compared to DiVE. This suggests that constraining the difference vectors to be zero imposes overly strong constraints, which may conflict with the objective of fine-tuning, i.e., adapting to the target data. We also found that it does not sufficiently preserve the geometric structure of the embeddings, as shown in Table 1². These findings suggest that allowing non-zero difference vectors, as DiVE does, provides greater flexibility during fine-tuning.

5.5 Ablation Study

We conducted an ablation study to analyze the effects of AVL and PVL on DiVE’s performance. Table 5 presents the results of applying AVL and PVL individually when using ImageNet as a target task. AVL significantly improved performance in OOD and zero-shot settings by 3.4 and 13.4 points, respectively. PVL also yielded improvements in these metrics, though to a lesser extent. This is because

²We used the same hyper-parameters as used by (Yu et al. 2024). While increasing the effect of the constraint by tuning the hyper-parameter improves the RSA correlation score and zero-shot performance, we empirically observed that it severely degrades ID performance, indicating that the fine-tuning process is impeded.

Method	ID	OOD avg.	ZS avg.	RSA
Cosine	82.4	62.3	61.7	0.949
DiVE	82.5	63.2	63.7	0.981

Table 7: Comparison between constraints on cosine similarities and difference vectors in DiVE.

PVL focuses solely on constraining the difference vectors for paired image-caption data samples without ensuring consistency across all difference vectors. However, when PVL is combined with AVL, it preserves the local geometric structure, thus complementing AVL and enhancing performance. These results indicate the effectiveness of complementary AVL and PVL, consistent with the results shown in Table 1.

We also investigated the effect of α in Eq. (6). Table 6 lists the results for various α in AVL. When $\alpha = 0$, both the OOD and zero-shot performance remained sub-optimal. This may be because the average vector m varies across batches, hindering the effective preservation of the geometric structure. Increasing α leads to improved performance. In particular, using a large α (e.g., 0.9 or 0.99) achieved high OOD and zero-shot performance. These results suggest that AVL enhances OOD and zero-shot performance by providing a stable average vector across batches.

6 Detailed Analysis

This section provides detailed analyses of DiVE. We conducted experiments on ImageNet as a target task.

Comparison to Cosine Similarity-based Constraint. As described in Sec. 2, CyCLIP (Goel et al. 2022) uses the cosine similarity-based constraint to learn a desired geometric structure. In contrast, DiVE uses the difference vector-based constraint. To compare these approaches, we modified DiVE to constrain the cosine similarities between embeddings instead of its original vector-based constraint. Specifically, we replaced the difference vectors, $u(\cdot)$ and $v(\cdot)$, in Eqs. (6), (7), and (8) with $u^{\cos}(\cdot)$ and $v^{\cos}(\cdot)$. They are defined as

$$u^{\cos}(\mathbf{x}_j^{\text{ref}}) = f_{\theta_{\text{ft}}}(\mathbf{x}_j^{\text{ref}}) \cdot f_{\theta_{\text{pre}}}(\mathbf{x}_j^{\text{ref}}), \quad (10)$$

$$v^{\cos}(\mathbf{t}_j^{\text{ref}}) = g_{\phi_{\text{ft}}}(\mathbf{t}_j^{\text{ref}}) \cdot g_{\phi_{\text{pre}}}(\mathbf{t}_j^{\text{ref}}). \quad (11)$$

Table 7 shows the comparison results, where ‘‘Cosine’’ indicates the use of $u^{\cos}(\cdot)$ and $v^{\cos}(\cdot)$. Constraining cosine similarities, as done with CyCLIP, resulted in a lower RSA correlation score compared with constraining the difference vectors. It also performed poorly across all evaluation metrics. This is likely because cosine similarity captures angular closeness on hyperspherical embeddings but not the direction of their difference. This result highlights the effectiveness of DiVE, which constrains the difference vectors.

Results using ViT-L/14. In Sec. 5, we used CLIP ViT-B/16 as the base model. To evaluate the generality across architectures, we report results using CLIP ViT-L/14 pre-trained on LAION-400M in Table 8. For methods that use the reference dataset (CC3M), i.e., FLYP + replay, SnD, and DiVE, we decreased B' to 256 due to GPU memory constraints. We observe that similar trends hold for ViT-L/14 as with ViT-B/16, i.e., DiVE achieved strong results across the three

Method	ID	OOD avg.	ZS avg.
Pre-training	75.2	70.9	72.7
FLYP	86.0	71.5	55.4
FLYP + replay	85.8	72.6	65.9
SnD	86.0	73.2	61.2
DiVE	86.1	74.5	70.1

Table 8: Experimental results on ImageNet using CLIP ViT-L/14. Bold indicates best result among fine-tuned methods, i.e., excluding pre-trained model.

Reference dataset	Size	ID	OOD avg.	ZS avg.
None (FLYP)	0	82.2	59.5	49.5
Flickr8K	8K	82.2	60.6	55.2
COCO Captions	118K	82.4	62.3	62.2
CC3M (default)	3M	82.5	63.2	63.7

Table 9: Results of DiVE with various reference datasets.

metrics. This indicates that the effectiveness of DiVE is not tied to a specific architecture and scales well to larger models, such as ViT-L/14.

Reference Datasets. We evaluated the impact of the choice of the reference dataset. Table 9 presents the results of DiVE with Flickr8K, COCO Captions, and CC3M as the reference dataset. We confirm that all variations consistently outperformed the baseline, which does not use a reference dataset (i.e., FLYP). We also observed that performance improved as the dataset size increased. This trend may be because larger datasets cover a wider region in the embedding space. Despite this trend, it is worth noting that, even when using COCO Captions, DiVE improved both OOD and zero-shot performance over ARF (relying on CC3M).

7 Conclusion

We proposed Difference Vector Equalization (DiVE), a robust fine-tuning method for contrastive pre-trained vision-language models. DiVE aims to preserve the geometric structure of the embeddings learned by pre-training, enabling fine-tuning without compromising the generalization ability. It thus constrains the difference vectors caused by fine-tuning to be equal across various data samples. We introduced two losses: average vector loss (AVL) and pairwise vector loss (PVL). AVL preserves the geometric structure globally by constraining difference vectors to be equal to their weighted average. PVL preserves the geometric structure locally by ensuring a consistent image-caption alignment. Extensive experiments demonstrated that DiVE significantly improves the performance of baseline methods by preserving the geometric structure. We believe that DiVE is not only a practical method but also opens a new direction in fine-tuning vision-language models by offering strong evidence for the importance of preserving the geometric structure of the embeddings. As future work, we aim to deepen the theoretical understanding of DiVE and clarify the principles underlying its effectiveness.

References

- Barbu, A.; Mayo, D.; Alverio, J.; Luo, W.; Wang, C.; Gutfreund, D.; Tenenbaum, J.; and Katz, B. 2019. ObjectNet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Beery, S.; Cole, E.; and Gjoka, A. 2020. The iWildCam 2020 Competition Dataset. *arXiv preprint, arXiv:2004.10340*.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101 – Mining Discriminative Components with Random Forests. In *European Conference on Computer Vision (ECCV)*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *International Conference on Machine Learning (ICML)*.
- Christie, G.; Fendley, N.; Wilson, J.; and Mukherjee, R. 2018. Functional Map of the World. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing Textures in the Wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Housley, N. 2021. An Image is Worth 16 × 16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations (ICLR)*.
- Dwivedi, K.; and Roig, G. 2019. Representation Similarity Analysis for Efficient Task Taxonomy & Transfer Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2004. Learning Generative Visual Models from Few Training Examples: An Incremental Bayesian Approach Tested on 101 Object Categories. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Gao, T.; Yao, X.; and Chen, D. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Goel, S.; Bansal, H.; Bhatia, S.; Rossi, R. A.; Vinay, V.; and Grover, A. 2022. CyCLIP: Cyclic Contrastive Language-Image Pretraining. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Goyal, S.; Kumar, A.; Garg, S.; Kolter, Z.; and Raghunathan, A. 2023. Finetune Like You Pretrain: Improved Finetuning of Zero-Shot Vision Models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Han, J.; Lin, Z.; Sun, Z.; Gao, Y.; Yan, K.; Ding, S.; Gao, Y.; and Xia, G. 2024. Anchor-based Robust Finetuning of Vision-Language Models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 2019. EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7): 2217–2226.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; Song, D.; Steinhardt, J.; and Gilmer, J. 2021a. The Many Faces of Robustness: A Critical Analysis of Out-of-Distribution Generalization. *International Conference on Computer Vision (ICCV)*.
- Hendrycks, D.; Zhao, K.; Basart, S.; Steinhardt, J.; and Song, D. 2021b. Natural Adversarial Examples. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ilharco, G.; Wortsman, M.; Gadre, S. Y.; Song, S.; Hajishirzi, H.; Kornblith, S.; Farhadi, A.; and Schmidt, L. 2022. Patching open-vocabulary models by interpolating weights. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.; Li, Z.; and Duerig, T. 2021. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In *International Conference on Machine Learning (ICML)*.
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. MaPLe: Multi-Modal Prompt Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Koh, P. W.; Sagawa, S.; Marklund, H.; Xie, S. M.; Zhang, M.; Balsubramani, A.; Hu, W.; Yasunaga, M.; Phillips, R. L.; Gao, I.; Lee, T.; David, E.; Stavness, I.; Guo, W.; Earnshaw, B.; Haque, I.; Beery, S. M.; Leskovec, J.; Kundaje, A.; Pierson, E.; Levine, S.; Finn, C.; and Liang, P. 2021. WILDS: A Benchmark of in-the-Wild Distribution Shifts. In *International Conference on Machine Learning (ICML)*.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3D Object Representations for Fine-Grained Categorization. In *IEEE International Conference on Computer Vision (ICCV) Workshops*.
- Kriegeskorte, N.; Mur, M.; and Bandettini, P. A. 2008. Representational similarity analysis - connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2.
- Kumar, A.; Raghunathan, A.; Jones, R.; Ma, T.; and Liang, P. 2022. Fine-Tuning can Distort Pretrained Features and Underperform Out-of-Distribution. In *International Conference on Learning Representations (ICLR)*.
- Lee, J. M. 2012. *Introduction to Smooth Manifolds*. Springer New York, NY.
- Lin, T.; Maire, M.; Belongie, S. J.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In *European Conference on Computer Vision (ECCV)*.

- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations (ICLR)*.
- Maji, S.; Kannala, J.; Rahtu, E.; Blaschko, M.; and Vedaldi, A. 2013. Fine-Grained Visual Classification of Aircraft. *arXiv preprint, arXiv:1306.5151*.
- Nilsback, M.; and Zisserman, A. 2008. Automated Flower Classification over a Large Number of Classes. In *Indian Conference on Computer Vision, Graphics & Image Processing*.
- Oh, C.; So, J.; Byun, H.; Lim, Y.; Shin, M.; Jeon, J.; and Song, K. 2023. Geodesic Multi-Modal Mixup for Robust Fine-Tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Parkhi, O. M.; Vedaldi, A.; Zisserman, A.; and Jawahar, C. V. 2012. Cats and Dogs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning (ICML)*.
- Rajaei, S.; and Pilehvar, M. T. 2021. How Does Fine-tuning Affect the Geometry of Embedding Space: A Case Study on Isotropy. In *Findings of the Association for Computational Linguistics: Empirical Methods in Natural Language Processing (EMNLP)*.
- Rashtchian, C.; Young, P.; Hodosh, M.; and Hockenmaier, J. 2010. Collecting Image Annotations Using Amazon’s Mechanical Turk. In *Annual Meeting of the Association for Computational Linguistics (ACL) HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*.
- Recht, B.; Roelofs, R.; Schmidt, L.; and Shankar, V. 2019. Do ImageNet Classifiers Generalize to ImageNet? In *International Conference on Machine Learning (ICML)*.
- Rolnick, D.; Ahuja, A.; Schwarz, J.; Lillicrap, T.; and Wayne, G. 2019. Experience Replay for Continual Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Schuhmann, C.; Vencu, R.; Beaumont, R.; Kaczmarczyk, R.; Mullis, C.; Jitsev, J.; and Komatsuzaki, A. 2021. LAION-400M: Open Dataset of CLIP-Filtered 400 Million Image-Text Pairs. In *Neural Information Processing Systems (NeurIPS) Data-Centric AI Workshop*.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Shu, Y.; Guo, X.; Wu, J.; Wang, X.; Wang, J.; and Long, M. 2023. CLIPood: Generalizing CLIP to Out-of-Distributions. In *International Conference on Machine Learning (ICML)*.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *arXiv preprint, arXiv:1212.0402*.
- Wang, H.; Ge, S.; Lipton, Z.; and Xing, E. P. 2019. Learning Robust Global Representations by Penalizing Local Predictive Power. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Wang, L.; Zhang, X.; Su, H.; and Zhu, J. 2024. A Comprehensive Survey of Continual Learning: Theory, Method and Application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8): 5362–5383.
- Wang, T.; and Isola, P. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In *International Conference on Machine Learning (ICML)*.
- Wortsman, M.; Ilharco, G.; Kim, J. W.; Li, M.; Kornblith, S.; Roelofs, R.; Gontijo-Lopes, R.; Hajishirzi, H.; Farhadi, A.; Namkoong, H.; and Schmidt, L. 2022. Robust fine-tuning of zero-shot models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xiao, J.; Hays, J.; Ehinger, K.; Oliva, A.; and Torralba, A. 2010. SUN Database: Large-scale Scene Recognition from Abbey to Zoo. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xu, J.; and Durrett, G. 2018. Spherical Latent Spaces for Stable Variational Autoencoders. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Yamaguchi, S.; Feng, D.; Kanai, S.; Adachi, K.; and Chijiwa, D. 2025. Post-pre-training for Modality Alignment in Vision-Language Foundation Models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yu, Y.; Huang, C.; Chen, J.; Chang, K.; Lai, Y.; Yang, F.; and Wang, Y. F. 2024. Select and distill: Selective dual-teacher knowledge transfer for continual learning on vision-language models. In *European Conference on Computer Vision (ECCV)*.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022a. Conditional Prompt Learning for Vision-Language Models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022b. Learning to Prompt for Vision-Language Models. *International Journal of Computer Vision (IJCV)*.